

SYSTEM AND METHOD FOR INTERCONNECTING MULTIPLE VIRTUAL PRIVATE NETWORKS

5 This application claims priority to the following provisional patent applications, which are incorporated herein by reference in their entireties:

(1) Provisional Application Serial No. 60/151,563, titled "Method & Apparatus For a Globalized Automotive Network & Exchange," filed on August 31, 1999, and having reference no. 99,532 (479.83581).

10 BACKGROUND OF THE INVENTION

Field of the Invention

20 The present invention relates to virtual private networks. More particularly, the present invention relates to virtual private networks wherein in each virtual private network, multiple service providers can be utilized by the trading partners of the virtual private network. The end-to-end service quality of the connection within the virtual private network is guaranteed to meet minimum requirements. The end-to-end service quality encompasses numerous factors including: network services; interoperability; performance; reliability; disaster recovery and business continuity; security; customer care; and trouble handling. The system and method of the present invention is directed to the interconnection of multiple virtual private networks each having multiple service providers. Furthermore the present invention encompasses a system and method for interconnecting multiple interconnect providers, such as exchange points, exchange networks, direct connect or transit service providers, between the multiple virtual private networks. Finally, the present invention employs an end-to-end overseer across the multiple virtual private networks.

Description of the Related Art

30 Early in 1994, the automotive industry recognized the need for global network services that would support more new demanding automotive business applications. The purpose of this network service was to simplify complex, redundant, outdated connection methods while minimizing costs and ensuring the management, security, reliability, and

performance essential to the automotive industry. Transport Control Protocol/Internet Protocol (TCP/IP) was endorsed as the standard suite for electronic data communications.

Ultimately in 1995, the industry formed a Telecommunications Project Team to oversee the design and development of a common global communication infrastructure supporting automotive industry application initiatives (later called the Automotive Network eXchange (ANX) Implementation Task Force). The Task Force, in June 1997, published the initial results of the technical design process for this new network service, called the Automotive Network eXchange (ANX), in "ANX Release 1 Draft Document Publication" (TEL-2 01.00). This reference is incorporated herein by reference in its entirety. The TEL-2 specification undergoes constant updating and correction.

The ANX system is a business-to-business communications infrastructure that provides a uniform, secured link between trading partners, such as manufacturers and suppliers, in the automotive industry. The ANX is a subscription-based network composed of Certified Service Providers (CSP). CSPs are providers of IP network service that have satisfied certain service end-to-end quality. CASPs are certificate authority service providers. The Certified Exchange Point Operator (CEPO) provides services to interconnect CSPs. CEPOs also must satisfy certain end-to-end service quality requirements.

Trading Partners (TP) are registered end users, or subscribers, of the ANX system such as automotive parts manufacturers, suppliers, original equipment manufacturers, and car manufacturers. The ANX system allows TPs to communicate, exchange information, and transact business with other TPs over the ANX network. The TP may utilize any TCP/IP-compliant application program to exchange information with other TPs. The registered TP selects the TPs with which it wants to communicate and thereafter may gain access to and receive communications from those selected TPs. As a result, the ANX system allows each TP to develop its own virtual private network with its customers and vendors.

The ANX system significantly reduces the complexity of connecting to multiple trading partners. Since there are diverse communication protocols for the trading partners, separate links are required to access each trading partner.

By having a single private network operated under a uniform protocol, interconnectivity between various trading partners is substantially simplified. In addition, ANX offers improved end-to-end service quality. For example, if an auto manufacturer needs to place with its parts supplier an order for car seats, the manufacturer may submit
5 over the ANX system its confidential CAD drawings directly to the supplier. The manufacturer may also fill out the order form that the supplier may have for filling orders and timely submit over the ANX system due to its high reliability and performance.

The CSP and the CEPO must satisfy certain performance and security requirements in order to be certified under the ANX. The certification process is
10 disclosed in ANX Release 1 Document Publication (TEL-2 02.00), which is incorporated herein by reference in its entirety.

The ANX VPN permits the use of a plurality of different IPsec devices. By virtue of the TEL-2 specification and the certification process all of the designated IPsec device are guaranteed to communicate with one another across the ANX VPN.

15 While the ANX was originated out of the need to interconnect automotive related companies, it is not limited to that industry. Any company/industry may become a TP, e.g. an aerospace company, a healthcare company, etc. ANX has become known as the Advanced Network eXchange.

With the advent of the Internet, global communication has become a reality.
20 While the Internet works well for non-mission critical applications, such as transmitting and receiving e-mail and hosting websites, it has some drawbacks for business-to-business commerce and communication that require stringent end-to-end service quality. Quality concerns are in the area of end-to-end service quality as explained previously.

For example, when two companies want to communicate over the Internet, the lag
25 between the systems at each company will be different virtually every time. The connection each has through their service provider, i.e. 14.4K, 28.8K, 56K, ISDN, DSL, T1, etc., plus the number of servers through which the connection is directed contribute to the resulting time lag between the two companies. Depending upon the type of information transmitted, the two parties may require a maximum acceptable time lag.
30 Due to the nature of the Internet, it cannot guarantee such a maximum time lag.

Furthermore, the two companies may desire that service assistance be available at certain times or 24 hours a day. The Internet has no such guarantees for help availability in a multi-provider environment. Such a lack of guaranteed bandwidth, latency and reliability are major impediments to business-to-business commerce and communication over the Internet.

In recent years the number of electronic viruses and hacker attacks has increased dramatically. A company considering conducting business-to-business commerce over the Internet runs the risk of making their intranet vulnerable to such viruses and attacks with the potential related loss of data.

In order to address the security issue, some companies have developed virtual private networks (VPNs). Secure VPNs permit a company to communicate with any other entity on the network without the risk of increased vulnerability to viruses and hackers. However, while VPNs can connect to other VPNs over the Internet by providing authentication, access control, confidentiality and data integrity, there is still no way the end-to-end quality of the connection can be guaranteed to meet a required set of minimum standards in a multi-provider setting.

A secure VPN is a communication network that is secured with encryption and authentication. Secure VPNs are based on multiple technologies, for example IPSec, tunneling, certification and shared secret authentication. IPSec is the security standard established by the Internet Engineering task Force (IETF). Tunneling permits private networks to cross the Internet using unregistered IP addresses.

SUMMARY OF THE INVENTION

From the foregoing, it is desirable to provide a system and method for interconnecting multiple VPNs each using multiple service providers while offering a minimum standard of end-to-end service quality.

The system and method of the present invention utilizes an overseer that defines the service quality, continually qualifies service providers as meeting that service quality, and resolves end-to-end issues across multiple interconnected virtual private networks, such as the ANX. When connecting multiple virtual private networks according to the system and method of the present invention multiple interconnect providers are

interconnected, and the manner in which these interconnect providers are interconnected so that the quality and reliability standards is met are another aspect of the present invention.

5 Certification of IPSec devices permits interoperability for encryption, integrity and authentication across the product of all IPSec vendors. When two subscriber companies both use certified IPSec equipment then they can provide each other with controlled access to each other's networks.

10 Based on the foregoing, an object of the present invention is to provide a system and method of interconnecting multiple VPNs each using multiple service providers while offering a minimum standard of end-to-end connection quality and reliability.

Another object of the present invention is to provide a system and method of interconnecting multiple VPNs having an overseer that resolves end-to-end issues across multiple virtual private networks.

15 Still another object of the present invention is to provide a system and method of connecting multiple virtual private networks in which multiple interconnect providers are interconnected so that the end-to-end service quality is met.

DETAILED DESCRIPTION OF THE DRAWINGS

The foregoing and other attributes of the present invention will be described with respect to the following drawings in which:

20

Fig. 1 is a block diagram of two interconnected virtual private networks according to the present invention;

25 **Fig. 2** is a configuration of governance and management of separate virtual private networks;

Fig. 3 is a configuration of governance and management of interconnected virtual private networks according to the present invention;

Fig. 4 is an interconnection configuration for governance of multiple interconnected virtual private networks according to the present invention;

5 **Fig. 5** is a flow chart showing contractual obligations according to the present invention;

Fig. 6 is a diagram illustrating end-to-end latency in a virtual private network having multiple service providers;

10 **Fig. 7** is a diagram illustrating end-to-end availability in a virtual private network having multiple service providers;

Fig. 8 is a diagram illustrating trouble handling in a virtual private network having multiple service providers;

15 **Fig. 9** is a diagram illustrating an accountability model for a single virtual private network having multiple service providers;

Fig. 10 is a diagram illustrating an accountability model for multiple virtual private networks having multiple service providers according to the present invention;

Fig. 11 is a diagram illustrating end-to-end interconnection of two virtual private networks according to the present invention;

25 **Fig. 12** is a diagram illustrating a trouble escalation model for interconnection of two virtual private networks according to the present invention;

Fig. 13 is a diagram illustrating a multiple virtual private network fee model for interconnection of two virtual private networks according to the present invention; is a

diagram illustrating interconnection of two virtual private networks using a multiple transit certified service providers according to the present invention;

5 **Fig. 14** is a diagram illustrating interconnection of two virtual private networks using a single transit certified service provider according to the present invention;

Fig. 15 is a diagram illustrating interconnection of two virtual private networks using a multiple transit certified service providers according to the present invention;

10 **Figs. 16** is a diagram illustrating interconnection of multiple virtual private networks using a multiple transit certified service providers, where no single transit certified service provider connects all of the virtual private networks according to the present invention; and

15 **Figs. 17a - c** are alternative configurations for interconnecting multiple virtual private networks according to the present invention.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENTS

Fig. 1 shows a block diagram of two interconnected virtual private networks 20 and 22. The present system and method of the interconnecting multiple virtual private networks is not intended to be limited to only these types of networks and has applicability to a wide variety of virtual private networks.

Each virtual private network 20 and 22 is shown having a trading partner (TP) 24 and 26, respectively. While Fig. 1 shows only one TP 24 and 26 for each virtual private network, there can in fact be hundred or thousands of such TPs for each virtual private network. Fig. 1 is intended to define the end-to-end service quality concept, and for such a purpose, only one TP 24 and 26 is need for each virtual private network 20 and 22.

The end-to-end service quality, provided by the present system and method of interconnecting multiple virtual private networks, cannot be achieved by simply interconnecting two virtual private networks, such as 20 and 22, with a wire. The end-to-end service quality incorporates a user-centric philosophy, where the user is the TP or subscriber. The user is guaranteed a minimum level of service encompassing factors that include: network services; interoperability; performance; reliability; disaster recovery and business continuity; security; customer care; and trouble handling. Simply connecting the two virtual private networks 20 and 22 with a wire will not achieve the minimum satisfactory levels for these factors.

To achieve such minimum levels of satisfactory performance for these factors the system and method must include a way to resolve disputes between the two virtual private networks. Referring to Fig. 2, each VPN 20 and 22 is shown as having its own governance, program management, coopetition policy, contracts, service assurance, and service description. While each virtual private network can operate with a successful level of end-to-end service quality when each VPN is not interconnected to another VPN, the governance, program management, coopetition policy, contracts, service assurance, and service description may need to be revised when interconnecting two or more VPNs in order to maintain the end-to-end service quality. It will be appreciated that at the very

least the interconnection of at least two VPNs adds at least one additional level of complexity with regard to service between the VPNs.

One resolution is shown in Fig. 3, in which each VPN 20 and 22 maintain their own governance, but the program management, cooperation policy, contracts, service assurance, and service description for the two VPNs 20 and 22 are unified. Such unification means that where the parameters for the program management, cooperation policy, contracts, service assurance, and service description of the two VPNs 20 and 22 are different, the parameter used in one of the networks is chosen as the acceptable minimum standard or a compromise parameter different from the parameter used in each or the VPNs is agreed upon. It is possible that the parameters for communication within each VPN need not change, while the new parameters are used only when communicating between VPNs. Fig. 3 further shows that the system and method contemplate connecting more than two VPNs.

One configuration for governance of multiple interconnected VPNs is shown in Fig. 4. In this scenario each VPN has its own program overseer (POVER) 30, and a global, or multiple virtual private network, overseer 32 is provided to resolve issues between the POVERs 30. Arrows are shown between the POVERs 30 indicating that the POVERs 30 are free to resolve their issues without requiring the GOVER 32. The GOVER is called on when direct POVER-to-POVER resolution fails. Each of the POVERs 30 governs one of the regional VPNs, while the GOVER 32 oversees the interconnection of the VPNs.

The GOVER is responsible for end-to-end quality assurance, and in particular acts as an inter-VPN interconnection certifier. The GOVER certifies interconnection facilities, and certifies a global CASP-CASP trust model. The GOVER also is an inter-VPN arbitrator that steps in when POVERs cannot resolve trouble between them.

Since the VPNs are used to running their networks in isolation, the interconnection of multiple VPNs has unique issues such as resolving trouble and conflicts between the VPNs and maintenance of minimum end-to-end service quality across the multiple programs. Since the system and method of the present invention are directed to providing specific end-to-end service quality, it must be possible for TPs to

quantify the end-to-end service quality levels, and these service quality levels must be sufficient to allow applications to work across the multiple VPNs. Therefore, a high level of metric compatibility and measurement techniques are required.

In the ANX type VPN each TP, CSP and CEP must meet specified criteria to become certified and to maintain that certification. The certification provides the TPs or subscribers with confidence that the level of transport and security will meet their business needs. The ANX type VPN utilizes multiple CSPs. On one level it is easier to run a VPN where all TPs are required to use a single CSP. The use of multiple CSPs in the ANX type VPN fosters competition between the CSPs and allows the VPN to reach TPs that may not be serviced by a single CSP. The implementation of multiple CSPs, however, brings with it the drawback of insuring that the CSPs can talk to one another. Whether the connection from one TP to another TP within the same VPN is through a single CSP or two CSPs should be invisible to the TPs. The TPs need never know when one or more CSPs are used for any particular connection. The certification process ensures that the TPs use one of the certified IPsec devices at their premises, and that the CSPs will utilize certified equipment and meet certain metrics so as to achieve the end-to-end service quality guaranteed to the TPs. In this manner, the multiple CSPs will be able to communicate with one another. The CSPs must meet business criteria, technical metrics, ongoing monitoring, trouble-handling criteria, routing registry criteria, and domain name registry criteria to achieve and maintain certification.

Fig. 5 shows the contractual obligations of the members of an ANX-type VPN. The TPs contract with the VPN, as denoted in Fig. 5 by the arrows to the overseer, and contract with one of the multiple CSPs. The CSPs contract with the VPN and with the CEPO. The CEPO contracts with the VPN. Each entity is responsible for the services that that entity provides.

The technical metrics for achieving end-to-end service quality in the ANX-type network include among other metrics, latency and availability. Fig. 6 illustrates the end-to-end latency within the ANX network. The TP1 router is connected to ANX CSP₁, which in turn is connected to ANX CEPO. TP2 router is connected to ANX CSP₂, which is connected to ANX CEPO. The packet latency from each router

and 66 through the corresponding CSP is 125 msec. The latency through the ANX CEPO is on the order of microseconds. The total packet latency through the network is therefore only slightly more than 250 msec.

Fig. 7 illustrates the end-to-end availability metric. The Access network between the TP1 router 60 and the ANX CSP₁ 62 is permitted to be unavailable 43.80 hours/year. The ANX CSP₁ 62 may only be unavailable 2.63 hrs./year. The trunk 65 between the ANX CSP₁ 62 and the ANX CEPO may only be unavailable 1.76 hrs./year. The ANX CEPO may only be unavailable 0.44 hours/year. The foregoing availabilities yield a total of 99.895% availability or 9.22 hours per year downtime.

The outline for how trouble is handled within the ANX-type VPN is shown in Fig. 8. There are effectively five layers of trouble handling. At the first level trouble between TPs is handled directly between the two TPs. Similarly, issues between the TPs and the CSPs are handled between the two parties. CSPs and the CEPOs also resolve their troubles between the troubled parties. A network overseer is provided to handle troubles that cannot be handled in the foregoing scenarios. The overseer can take complaints from the TPS, the CSPs, and the CEPOs.

A key to providing predictable end-to-end service quality is that the TPs must know the level of service they receive. To this end four service provider accountability levels exist. First, service providers, both interconnect providers and CSPs, must timely fix infrequent service provider troubles. Second, there must be end-to-end service provider cooperation to handle any troubles. Third, recourse must be provided to resolve disputes in the event of disagreement between CSPs and/or interconnect providers. Fourth, recourse must be provided to resolve continued non-compliance with the end-to-end service quality.

Referring to Figs. 9 and 10, charts for single VPN and interconnected VPNs are shown, respectively. In Fig. 9, the CSPs 70, CEPOs 72 and CASPs 74 are accountable to the POVER 76. The POVER 76 is accountable to the body 78 representing the TPs. The body 78 is accountable a regional/national arbitration body 80. Where multiple VPNs are interconnected in Fig. 10, the CSPs 70, the CEPOs 72, and CASPs 74 are accountable to the POVERs 76. The POVERs 76 are accountable to a GOVER 77, which in turn is

accountable to the body 78. The body 78, instead of being accountable to the regional/national arbitration body 80, is accountable to an international arbitration body 82.

The GOVER/POVER model is but one way to oversee ensuring of the end-to-end service quality and metric compatibility. How the ANX-type networks are connected will be discussed below. In this context there must be five key types of end-to-end technology compatibility: 1 network interconnection that ensures a trading partner on one VPN can reach any trading partner on the other VPN; 2 routing compatibility that ensures any trading partner on one VPN can logically reach any TP on the other VPN; 3 naming compatibility, e.g. so the web names or e-mail names of any trading partner on one VPN can be resolved to an address that is routable over the two VPNs; 4 IPsec compatibility; and 5 digital security certificate compatibility across multiple VPNs. While Figs. 9 and 10 refer to regional/national VPNs and international arbitration, the VPNs need not be limited to a specific country or geographical area. Any ANX-type VPN, regardless of the location of its subscribers could be interconnected.

While Fig. 1 illustrated the interconnection of two VPNs 20 and 22, a significant element is missing. Fig. 11 shows two VPNs, that have multiple service providers, which are connected through an inter-program service provider, also called an interconnect provider. The Tel-2 specification is still used as the basic guide in determining the end-to-end service quality, however regional or particular VPN peculiarities, referred to as deltas, must be considered when establishing the interconnected end-to-end service quality standards.

Returning to the GOVER/POVER model for overseeing interconnected VPNs; Fig. 12 illustrates an end-to-end trouble escalation model. It is expected that CSPs will work together to resolve trouble before contacting a POVER. Similarly, the POVERs and/or the POVERS and the interconnect provider are expected to work together to resolve trouble before contacting the GOVER.

When expanding from a single VPN to interconnected VPNs the inherent costs of running the system naturally increase. How such costs are distributed is an important part of the system. As shown in Fig. 13, the POVERs 100 pay fees to the GOVER to

offset the costs of maintaining the GOVER. The VPNs having multiple service providers in turn pay fees to support the POVER. Furthermore the interconnect providers pay a certification fee to the GOVER for certification as a interconnect provider between VPNs.

5 There are multiple methods of interconnecting multiple VPNs with interconnect providers. As shown in Fig. 14, all the VPNs, having multiple service providers, can be interconnected using a single interconnect provider. Alternatively, all the VPNs can be interconnected by multiple interconnect providers, as shown in Fig. 15, thereby creating competition between the interconnect providers, just as there is competition between the
10 CSPs in a single xNX-type VPN. Finally, as shown in Fig. 16, where no suitable interconnect provider is available to connect all he VPNs having multiple service providers, multiple interconnect providers are used. These interconnect providers service different combinations of VPNs. In Fig. 16, interconnect provider 120 interconnects VPNs having multiple service providers 122, 124, and 126. Interconnect provider 130
15 interconnects VPNs having multiple service providers 132 and 126. As a result, a TP of VPN 132 must connect through both Interconnect provider 130 and Interconnect provider 120 to reach TPs of either VPN 122 or 124.

How the multiple VPNs interconnect will directly affect the resulting end-to-end service quality. Figs. 17a-c illustrate potential configurations of multiple VPNs. In Fig.
20 17a a first TP 200 connects to a first CSP 210. The CSP210 connects to a first exchange point 220. The TP 200, CSP 210, and the exchange point 220 are within a first VPN 240. A second TP 250 connects to a second CSP 260, which connects to a second exchange point 270. The TP 250, CSP 260 and exchange point 270 are within a second VPN 280. The two VPNs 240 and 280 are interconnected by an Interconnect provider 300, which is
25 connected to the exchange points 220 and 270.

In Fig. 17b TP 200, CSP 210, exchange point 220 and Interconnect provider 300 are connected in the same manner shown in Fig. 17a. While the second TP 250 is connected to the CSP 260, the exchange point 270 is not provided. Instead CSP 260 is shown as connecting directly to the Interconnect provider 300. This embodiment

encompasses the situation where the Interconnect provider 300 and CSP 260 are the same entity or are directly wired. Fig. 17c is similar to Fig. 16b, Except that the interconnect provider also acts as a CSP 320, and a third TP 310 is connected directly to the Interconnect provider 300/CSP 320.

5 As stated previously, while the end-to-end service quality is based upon the TEL-2 specification, the degree to which the TEL-2 specification needs to be modified to interconnect multiple VPNs depends upon the chosen complexity of the interconnection. An xNX-type VPN uses a maximum of two CSPs between any two TPS. A larger value, either three or four, is needed for multiple VPNs. The Interconnect provider will account
10 for one additional CSP, and for configuration set forth in Fig. 16, two Interconnect providers are employed in addition to the two CSPs yielding four CSPs.

 Having described several embodiments of the system and method for interconnecting multiple virtual private networks in accordance with the present invention, it is believed that other modifications, variations and changes will be suggested
15 to those skilled in the art in view of the description set forth above. It is therefore to be understood that all such variations, modifications and changes are believed to fall within the scope of the present invention as defined in the appended claims.